Traffic-Driven Spectrum Allocation in Heterogeneous Networks

Binnan Zhuang, Dongning Guo, and Michael L. Honig

arXiv:1408.6011v2 [cs.IT] 26 Mar 2015

***Abstract*—Next generation cellular networks will be heterogeneous with dense deployment of small cells in order to deliver high data rate per unit area. Traffic variations are more pronounced in a small cell, which in turn lead to more dynamic interference to other cells. It is crucial to adapt radio resource management to traffic conditions in such a heterogeneous network (HetNet). This paper studies the optimization of spectrum allocation in HetNets on a relatively slow timescale based on average traffic and channel conditions (typically over seconds or minutes). Specifically, in a cluster with n base transceiver stations (BTSs), the optimal partition of the spectrum into 2^n segments is determined, corresponding to all possible spectrum reuse patterns in the downlink. Each BTS's traffic is modeled using a queue with Poisson arrivals, the service rate of which is a linear function of the combined bandwidth of all assigned spectrum segments. With the system average packet sojourn time as the objective, a convex optimization problem is first formulated, where it is shown that the optimal allocation divides the spectrum into at most n segments. A second, refined model is then proposed to address queue interactions due to interference, where the corresponding optimal allocation problem admits an efficient suboptimal solution. Both allocation schemes attain the entire throughput region of a given network. Simulation results show the two schemes perform similarly in the heavy-traffic regime, in which case they significantly outperform both the orthogonal allocation and the full-frequency-reuse allocation. The refined allocation shows the best performance under all traffic conditions.**

I. Introduction

Spectrum management in current cellular networks includes two stages: First, the locations of base transceiver stations (BTSs) and spectrum (carrier) assignments are carefully planned offline. Once deployed, each BTS scheduler allocates time-frequency resource blocks to users on a fast timescale. Spectrum assignments in early networks have often been based on regular lattice frequency reuse patterns. In current 4G networks, spectrum assignments are based either on full or fractional frequency reuse (FFR). In FFR, a main portion of the spectrum is reused everywhere except at cell edge, and the remaining spectrum is divided for orthogonal reuse at cell edge (e.g., [2]).

For next generation networks, aggressive frequency reuse through dense deployment of (small) micro/pico cells is a major means for overcoming the shortage of spectrum resources [3]. Such a heterogenous network (HetNet) with overlapping cells of all sizes often operates in the interference limited regime. Because small cells lead to more pronounced traffic and interference variations, traditional static frequency reuse is not effective [4]. Neither are the semi-regular reuse patterns of dynamic FFR (e.g., [5]–[7]).

In this paper, we introduce a model for a HetNet with dynamic traffic, present two optimization-based spectrum allocation schemes, and demonstrate their effectiveness using simulation. The timescale of resource adaptation here is conceived to be relatively slow, e.g., once every a few seconds or minutes. This timescale is, on the one hand, fast enough for tracking the aggregate traffic variation, and, on the other hand, slow enough to allow joint optimization of many cells with a large number of user equipments (UEs). This is in contrast to most existing work, which considers resource allocation on the timescale of a frame, assuming instantaneous information exchange between cells (see, e.g., [5]–[9]).

Because the period of spectrum allocation is much slower than the channel coherence time, the channel conditions are accurately modeled using path loss and the statistics of small-scale fading. Moreover, any given frequency band is assumed be homogeneous, i.e., the utility of any one Hertz of spectrum assigned to a cell depends only on the corresponding reuse pattern, i.e., the subset of cells that share it. Specifically, in an n-cell HetNet, there are exactly 2^n distinct reuse patterns. Indeed, the slow timescale allocation is fundamentally equivalent to deciding on the bandwidths of all those reuse patterns. This is formulated as a convex program with optimality guarantee and relatively low computational complexity. This is in contrast to most related work in the literature, where the allocation problem is formulated as that of deciding, for each slice of the spectrum, which BTSs should use it. The latter problem in general is a discrete optimization problem that can be hard to solve without resorting to heuristic methods, and may have many local optima (see, e.g., [10]–[14]).

One important feature of this work is the assumption of stochastic packet arrivals to each cell. Recent studies [15], [16] point out that the usual backlogged traffic assumption exaggerates the inter-cell interference in the small-cell scenario. Another feature is the choice of delay performance as the objective for optimization. Such a quality of service (QoS) metric is more relevant for a HetNet with dynamic traffic than the frequently used sum rate and the outage probability. Although there is a large body of literature on physical layer resource allocation, few papers use network layer QoS as the performance metric. Examples include the models in [17], [18], which correspond to full-frequency-reuse (or full-reuse) only, hence spectrum allocation is not considered. In particular,

The authors are with Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208, USA.

The work was presented in part at Globecom 2014 [1]. This work was supported in part by a gift from Futurewei Technologies and by the National Science Foundation under Grant No. CCF-1018578.

The manuscript includes materials to appear in IEEE Journal on Selected Areas in Communications as well as an appendix omitted in the journal paper.

resource allocation in [18] is performed in the time domain by iteratively updating a scheduling policy and BTS utilizations.

The resource allocation problem here is a joint physical layer and network layer optimization problem. To connect the spectrum (and power) resources in the physical layer to the QoS in the network layer, we use the service rates of the queues of all cells as the link. Specifically, the spectral efficiencies along with the bandwidths of all the reuse patterns allocated to a BTS determine the instantaneous service rate of the corresponding cell. The average packet delay in a given cell is in turn determined by the service rates and the packet arrival rates. Thus an optimization problem is formulated with the average packet sojourn time as the objective, the bandwidths of the reuse patterns as the desired variables, and the service rates as intermediate variables.

Two allocation schemes, referred to as the "conservative" scheme and the "refined" scheme, are obtained based on different service rate models described in Section II. The conservative scheme, discussed in Section III, assumes that a BTS's transmission rate over any spectrum segment is the worst-case rate under the corresponding reuse pattern, which is the achievable rate when all BTSs in the pattern are transmitting. In this case the queueing dynamics at one BTS does not depend on other BTSs' activities, so that its average delay is simply that of an M/M/1 queue. The resulting optimization problem is convex. An important finding is that the optimal allocation uses at most n out of all 2^n reuse patterns. That is, it suffices to divide the spectrum into n segments for allocation, rendering the solution practical.

In the second scheme, discussed in Section IV, the instantaneous service rate for one BTS's queue depends on whether interfering BTSs are transmitting. The analysis of n interactive queues is a long-standing open problem. To make progress, we develop an approximation for the average packet sojourn time as the objective for the optimization. Although the problem may not be convex in general, standard solvers for convex programming appear to perform remarkably well, and the number of active reuse patterns is very close to n.

Both allocation schemes are shown to achieve the entire throughput region of a given network. By simulating a cell cluster, their performance is compared with two other schemes, namely, full reuse and optimal orthogonal frequency reuse, in Section V. The refined allocation achieves the minimum delay under all circumstances. The proposed schemes show the largest improvement in the heavy-traffic regime, where their performance is very similar.

In independent work [19], Kuang et al. also considered the allocation of spectrum to an arbitrary set of reuse patterns, and incorporated user association as well. Their performance objective is sum logarithmic utility function, so traffic and queuing models are not included in the formulation. A gradient based numerical method is used to solve the optimization problem, which has similar computational complexity to the algorithms in this paper.

II. System Model

Suppose an operator uses one licensed frequency band of W Hz to carry all downlink transmissions in a HetNet of n cells. Denote the set of n BTSs as $N = \{1, \ldots, n\}$. On a slow timescale, the frequency resources are assumed to be homogeneous. In each period, the task of a central controller is to determine which part of the spectrum is allocated to each BTS. The problem is equivalent to deciding the bandwidth of all 2^n reuse patterns, denoted by a 2^n-tuple: $\boldsymbol{x} = (x_B)_{B \subset N}$, where $x_B \in [0, 1]$ is the fraction of spectrum shared by BTSs in set B. Clearly, $\sum_{B \subset N} x_B = 1$, and any efficient allocation would not use the empty reuse pattern, so that $x_\emptyset = 0$. An example allocation to 3 BTSs is shown in Fig. 1, where the spectrum is divided into 7 segments, corresponding to all $2^3 - 1 = 7$ nonempty reuse patterns. The spectrum allocated to a BTS can be noncontiguous, which may be implemented using orthogonal frequency division multiplexing (OFDM).

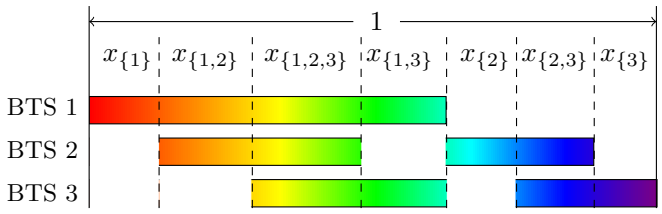

Fig. 1. An example of spectrum allocation to 3 BTSs. All $2^3 - 1 = 7$ possible nonempty reuse patterns are shown. The bandwidth of the spectrum shared by BTSs in set $A \subset \{1, 2, 3\}$ is denoted as x_A and the total bandwidth is 1.

In this paper, we assume that either a single UE is associated with each BTS or all UEs associated with a BTS are colocated, so that they can be regarded as a single user on the slow timescale. (This model can be generalized to allow any number of UEs at arbitrary locations, where the BTS-UE association is also to be optimized.) Let the aggregate traffic arriving at BTS i be an independent Poisson point process with rate λ_i packets per second. The length of each packet independently follows the identical exponential distribution with average packet length of L bits. Let packets intended for different UEs within a cell be processed according to the first-in-first-out (FIFO) criterion.[1]

For simplicity, it is assumed that when BTS i transmits, it employs all reuse patterns available to it and applies a flat power spectral density (PSD) p_i over the allocated spectrum. At any frequency, the instantaneous spectral efficiency achievable by BTS i depends on the set of active BTSs $A \subset N$ using that frequency. Let this spectral efficiency be denoted by $s_{i,A}$. Evidently, $s_{i,A} = 0$ if $i \notin A$. Moreover, the spectral efficiency decreases as more BTSs become active, i.e., $s_{i,A} \geq s_{i,B}$ if $i \in A \subset B$. On the slow timescale, the spectral efficiencies are either known *a priori* or can be computed or measured by the central controller. For later convenience, we convert the units of $s_{i,A}$ from bits/second/Hz to packets/second by normalizing it with L/W bits/packet/Hz. For concreteness in obtaining numerical results, we use Shannon's formula to obtain:

$$s_{i,A} = \frac{W\, 1(i \in A)}{L} \log_2 \left(1 + \frac{p_i}{I_{i,A}}\right) \quad \text{packets/second} \quad (1)$$

[1]The results in this paper apply to all queueing disciplines that are work-conserving, non-anticipating and non-preemptive as defined in [20].

where $1(i \in A) = 1$ if $i \in A$ and $1(i \in A) = 0$ otherwise, and $I_{i,A}$ is the total noise plus interference PSD from other BTSs in A to the UEs of cell i, depending on the transmit PSDs and path loss. The effect of small-scale fading can be included by considering the ergodic capacity in lieu of (1), which does not change the main results of this paper.

The cells form a system of n (interactive) queues. Two service rate models are conceivable: 1) In the so-called *conservative model*, BTS i transmits at rate $s_{i,B}$ over reuse pattern B, which is achievable regardless of the activities of other BTSs. The rate contributed by reuse pattern B is the product of the spectral efficiency and the bandwidth: $s_{i,B}x_B$. Hence the service rate of cell i is the sum rate of all reuse patterns, expressed as:

$$r_i = \sum_{B \subset N} s_{i,B} x_B \quad \text{packets/second.} \tag{2}$$

2) In the so-called *refined model*, every BTS adapts its rate to the instantaneous set of active BTSs, denoted as A. The rate contributed by reuse pattern B to cell i is $s_{i,B\cap A}x_B$. The service rate of cell i depends on A and is expressed as:

$$r_{i,A} = \sum_{B \subset N} s_{i,B\cap A} x_B \quad \text{packets/second.} \tag{3}$$

In general, the service rate in the refined model is higher, i.e., $r_{i,A} \geq r_i$, because $s_{i,B\cap A} \geq s_{i,B}$. In either case, the system is modeled as n continuous-time Markov chains (CTMCs).

In the next two sections, we develop two spectrum allocation schemes based on the preceding two service rate models, respectively. The objective is to minimize some queueing delay by optimizing the spectrum allocation $\boldsymbol{x}$.

III. A Conservative Spectrum Allocation Scheme

In this section, the service rate model (2) is assumed. This model is conservative in the sense that each BTS transmits at the worst-case rate that is achievable when all other BTSs are always interfering. This is equivalent to assuming that other cells' traffic is always backlogged. Under this model, the cells form n independent M/M/1 queues. The average packet sojourn time in cell i takes a simple form [21]:

$$t_i = \frac{1}{r_i - \lambda_i} \quad \text{seconds.} \tag{4}$$

Note that $1/(r-\lambda)$ is strictly convex in r on (λ, ∞).

A. The Optimization Problem

The spectrum allocation problem based on the conservative approximation (4) is formulated as:

$$\begin{aligned}
\underset{\boldsymbol{r},\boldsymbol{x}}{\text{minimize}} \quad & \frac{1}{\sum_{j=1}^n \lambda_j} \sum_{i=1}^n \frac{\lambda_i}{r_i - \lambda_i} && && \text{(P1a)}\\
\text{subject to} \quad & r_i = \sum_{B \subset N} s_{i,B} x_B, && \forall i \in N && \text{(P1b)}\\
& r_i > \lambda_i, && \forall i \in N && \text{(P1c)}\\
& x_B \geq 0, && \forall B \subset N && \text{(P1d)}\\
& \sum_{B \subset N} x_B = 1. && && \text{(P1e)}
\end{aligned}$$

The variables in the optimization are $\boldsymbol{r} = [r_1, \ldots, r_n]$ and $\boldsymbol{x}$. The objective (P1a) is the average packet delay of the entire network, where $\lambda_i / \sum_{j=1}^n \lambda_j$ is the fraction of total traffic in cell i. The constraints (P1c) guarantee the stability of all queues. Problem (P1) is a convex optimization problem because all constraints are linear and the objective is a linear combination of convex functions.

Since the objective is strictly convex and positive, the optimization problem (P1) has a unique global minimum when feasible [22]. Moreover, due to the special structure of the problem, the optimal allocation divides the spectrum into at most n segments.

Theorem 1: In the optimal solution of the n-BTS conservative spectrum allocation problem, at most n out of the 2^n reuse patterns are active, i.e., the optimal solution $\boldsymbol{x}$ statisfies

$$|\{B \mid x_B > 0,\ B \subset N\}| \leq n. \tag{5}$$

Proof: For every $B \subset N$, the spectral efficiency vector $\boldsymbol{s}_B = [s_{1,B}, \ldots, s_{n,B}]$ denotes a point in $\mathbb{R}^n$. According to (P1b) to (P1e), $\boldsymbol{r} \in \mathbb{R}_+^n$ is a convex combination of the 2^n points $(\boldsymbol{s}_B)_{B\subset N}$ with coefficients $(x_B)_{B\subset N}$, i.e., $\boldsymbol{r} = \sum_{B\subset N} \boldsymbol{s}_B x_B$. In other words, any $\boldsymbol{r}$ given by (P1b) is in the convex hull of $(\boldsymbol{s}_B)_{B\subset N}$. By Carathéodory's Theorem [23], $\boldsymbol{r}$ lies in a d-simplex with vertices in $(\boldsymbol{s}_B)_{B\subset N}$ and $d \leq n$, i.e., $\boldsymbol{r}$ can be written as a convex combination of $(x_B)_{B\subset N}$ with at most $n+1$ nonzero coefficients in $\boldsymbol{x}$. This holds for any $\boldsymbol{r}$ satisfying (P1b) to (P1e). Furthermore, the $\boldsymbol{r}^*$ corresponding to the optimal solution to (P1) must be Pareto optimal in terms of the rate allocation, i.e., one cannot find another spectrum allocation $\boldsymbol{x}$ that satisfies $r_i^* \leq \sum_{B\subset N} s_{i,B} x_B,\ \forall i \in N$ with at least one of the inequalities being strict. This is because any spectrum allocation that could increase the service rate at any BTS without decreasing the rates at other BTSs would also decrease the objective (P1a). Hence $\boldsymbol{r}^*$ cannot be an interior point of the d-simplex, and must lie on some m-face of the d-simplex with $m < d \leq n$. Therefore $\boldsymbol{r}^*$ can be written as a convex combination with $m+1 \leq n$ nonzero coefficients in $\boldsymbol{x}$. ■

Corollary 1: In the optimal solution of the n-BTS conservative spectrum allocation problem, for any subset $M \subset N$ of m BTSs, the spectrum exclusively used by those BTSs in M is divided into at most m segments:

$$|\{B \mid x_B > 0,\ B \subset M\}| \leq m. \tag{6}$$

Proof: If x_B for every $B \subset N$, $B \not\subset M$ is fixed (at its optimal value), then (P1) becomes an optimization problem over variables $(x_B)_{B\subset M}$. The service rates at BTSs not in M are fixed, and the service rates for the m BTSs in M are convex combinations of $(x_B)_{B\subset M}$ plus a constant vector in $\mathbb{R}_+^m$. The optimization problem reduces to the form of (P1) with only m BTSs. Hence Corollary 1 is proved using the same arguments used to prove Theorem 1. ■

B. An Efficient Algorithm

The optimization problem (P1) has $n + 2^n$ variables. The computational complexity is typically polynomial in 2^n using a standard convex optimization solver. The structure of the

optimal solution given by Theorem 1 suggests a more efficient solution, as we only need to determine the bandwidths of the n nonzero segments. The difficulty of course is to decide which n segments. Algorithm 1 solves the n-BTS spectrum allocation problem iteratively, the details of which are explained next.

Algorithm 1 The conservative spectrum allocation scheme

INPUT: λ_i and $s_{i,B}$ for all $i \in N$ and $B \subset N$.
OUTPUT: $(x_B)_{B \subset N}$.
Initialization: Find a feasible solution $(x'_B)_{B \subset N}$ by solving (P1) with constant objective. $\Omega \leftarrow \{B \mid x'_B > 0\}$, $\Omega' \leftarrow \emptyset$.
while $\Omega \not\subset \Omega'$ **do**
1. $\Omega' \leftarrow \Omega$;
2. Find $(x_B)_{B \subset N}$ by solving (P1) starting from $(x'_B)_{B \subset N}$ with additional constraints: $x_B = 0$, $\forall B \notin \Omega$;
3. Compute the partial derivatives of the objective function (P1a) with respect to each element in $(x_B)_{B \subset N}$, namely, $\Delta_{x_B} \leftarrow -\sum_{i \in B} \frac{\lambda_i s_{i,B}}{(r_i - \lambda_i)^2}$;
4. $\Omega \leftarrow \{B$ for the n smallest $\Delta_{x_B}\}$, $\Omega \leftarrow \Omega \cup \Omega'$, $(x'_B)_{B \subset N} \leftarrow (x_B)_{B \subset N}$.
end while

1) Initialization: To find a feasible point to start, we first solve a modified (P1) by replacing the objective function (P1a) with a constant. The problem can be transformed to a linear program in standard form, which can be solved using the simplex method [24]. Although the worst-case complexity of this method is $O(2^{2^n})$, for most practical problems the complexity is usually very low. According to the properties of basic feasible solution to a linear program [24], the solution (an initial point for (P1)) will have at most $n+1$ active reuse patterns, which are collected in the *set of candidate reuse patterns*, denoted as Ω.

2) Description of the Iterations: The main idea is similar to the delayed column generation algorithm for solving large-scale linear programs. The difference is the criterion used to select variables to be added to the set of candidate reuse patterns, Ω. In each iteration, the algorithm first finds the optimal solution within the candidate set Ω. Then, the partial derivatives with respect to x_B for all $B \subset N$ are calculated (including those not in Ω). The n reuse patterns with the smallest derivatives are added to the candidate set Ω, to be used in the next iteration. (The number of variables added to the candidate set may be fewer than n due to overlap.) The algorithm terminates when the candidate set ceases to grow. The proposed algorithm is guaranteed to converge to the global optimum, because otherwise it can always find at least one new reuse pattern to add to Ω. In the worst-case, the candidate set may eventually include all 2^n variables, but typically the algorithms terminates quickly.

3) Performance: Usually, the fewer reuse patterns the algorithm begins with, the faster it converges. Thus if the full-reuse allocation ($x_N = 1$) is feasible, it is a preferred initial point. At any rate, the initial solution has no more than $n+1$ active reuse patterns. Example plots of delay versus the number of iterations are shown in Fig. 2 with 7 cells and different traffic loads. In the simulation, Algorithm 1 starts with the full-reuse allocation and converges to the global optimum within a few iterations.

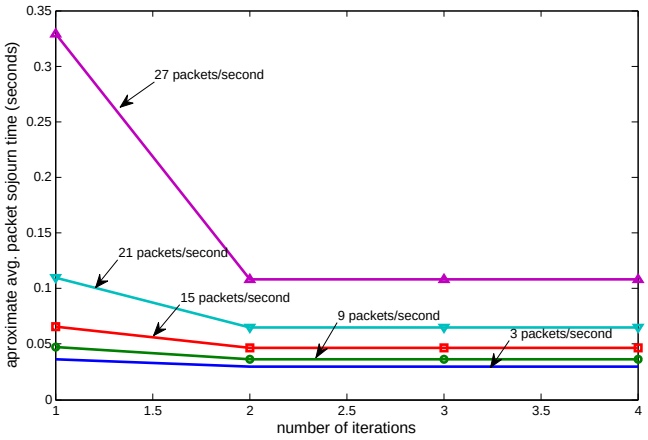

Fig. 2. Approximated average delay versus number of iterations using Algorithm 1 with different average packet arrival rates.

IV. A Refined Spectrum Allocation Scheme

In this section, we adopt the refined service rate model (3) and develop a corresponding spectrum allocation scheme. We assume each BTS adapts its transmit rate to the instantaneous interference level at its UEs. One way to implement this in practice without explicit knowledge of the interference levels is to use a hybrid automatic repeat request (HARQ) scheme.

The n cells form a system of n (interactive) queues, where the instantaneous service rate of each queue depends on which other queues are empty. Under such a *coupled-processors* model, it is challenging to express the queueing delay as a function of the spectrum allocation and the arrival rates. In the special case of two coupled queues, finding the joint steady-state distribution can be formulated as a Riemann-Hilbert problem [25]. The same two-dimensional Markov process has been studied in [26], where the steady-state distribution can be represented as an infinite series of product forms. Two coupled processors with generally distributed service times have been studied in [27], which shows the joint workload distribution can be determined by solving a boundary value problem. These results are difficult to use for numerical computation. Also, few results exist for more than two coupled queues.

In Section IV-A, we derive an approximation for the average packet sojourn time based on a modified queueing system, where the queue interactions are somewhat simplified. A spectrum allocation scheme is then developed in Section IV-B to minimize the delay approximation. In Section IV-C, we visit the best known upper and lower bounds on the actual delay. It is then shown in Section IV-D that the proposed approximation is between existing upper and lower bounds.

A. Delay Approximation

The original n-dimensional CTMC can be described as follows. Let $\boldsymbol{l} = [l_1, \ldots, l_n]$ be the state of the CTMC, where $l_i \geq 0$ is the number of packets in queue i (including the packet being served, if any). Let $A_{\boldsymbol{l}} = \{i \,|\, l_i > 0\}$ denote

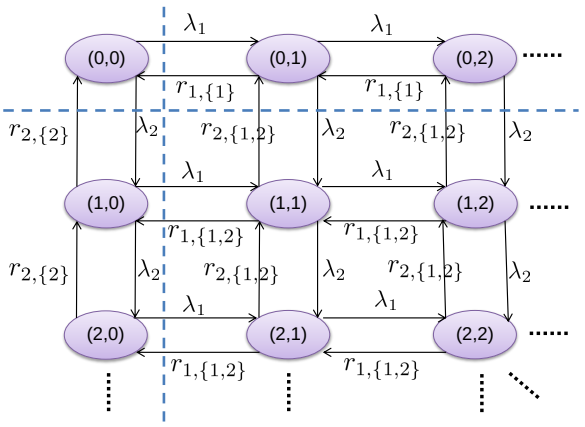

Fig. 3. The continuous time Markov Chain for the 2-BTS interactive queueing model.

the corresponding set of active BTSs. The transition rate from state $\boldsymbol{l}$ to state $\boldsymbol{l}' = [l'_1, \ldots, l'_n]$ is given by:

$$Q(\boldsymbol{l}, \boldsymbol{l}') = \begin{cases} \lambda_i, & \text{if } l_i + 1 = l'_i,\ l_j = l'_j\ \forall j \neq i \\ r_{i,A_{\boldsymbol{l}}}, & \text{if } l_i - 1 = l'_i,\ l_j = l'_j\ \forall j \neq i \\ -\sum\limits_{\boldsymbol{l}'':\boldsymbol{l}'' \neq \boldsymbol{l}} Q(\boldsymbol{l}, \boldsymbol{l}''), & \text{if } \boldsymbol{l} = \boldsymbol{l}' \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

where the service rate $r_{i,A}$ is defined in (3). As an example, the CTMC for a 2-BTS system is illustrated in Fig. 3.

To ease the analysis, we shall approximate the original CTMC by a modified CTMC with reduced memory. Because the service rate r_i depends only on the set of active queues, we group all the states corresponding to each active set A and refer to them as group A. According to (7), transitions occur between neighboring states where the length of a single queue increases or decreases by 1. Moreover, the CTMC makes a certain number of transitions between states within a group A before it jumps into a different group A'. Specifically, we make the following two (somewhat strong) assumptions about the modified CTMC:

1) When the modified CTMC transits from some state $\boldsymbol{l}$ in group A to some random state $\boldsymbol{l}'$ in a different group A', the new state is independent of $\boldsymbol{l}$. In other words, such a transition is memoryless;
2) When the modified CTMC transits from group A to a different group A', the probability of assuming any state in A' is proportional to the stationary distribution of the state in the modified CTMC.

In the following, we analyze the stationary distribution of the modified CTMC. By definition, each inter-group state transition of the CTMC is a renewal, because the state is chosen anew within the new group according to the steady-state distribution. In addition, the intra-group transitions within a group form an independent CTMC. An important observation is that, within a group, where the set of active BTSs is fixed, all the service rates are invariant with the queue lengths, so that the n queues become independent. The probability of a state $\boldsymbol{l}$ is thus decomposed as

$$p(\boldsymbol{l}) = p(A_{\boldsymbol{l}}) \prod_{i=1}^{n} p_i(l_i | A_{\boldsymbol{l}}) \tag{8}$$

where $p_i(l|A)$ is the probability that queue i has length l given that the state is in group A. Evidently, $p_i(0|A) = 0$ if $i \in A$, $p_i(0|A) = 1$ if $i \notin A$, and $p_i(l|A) = 0$ if $i \notin A$, $l > 0$. It suffices to determine the steady-state distribution of the groups $A \subset N$, and, for each group A, the steady-state distribution of the states in A.

1) The Intra-group CTMC: Consider the independent CTMC within any given group A. The CTMC can only exit group A when some queue length is equal to 1. For queue $i \in A$, the probabilities of states 1, 2, ... must satisfy the detailed balance equation:

$$p_i(l|A)\lambda_i = p_i(l+1|A) r_{i,A}, \quad l = 1, 2, \ldots. \tag{9}$$

As a result, for every $i \in A$,

$$p_i(l|A) = \left(1 - \frac{\lambda_i}{r_{i,A}}\right)\left(\frac{\lambda_i}{r_{i,A}}\right)^{l-1}, \quad l = 1, 2, \ldots. \tag{10}$$

It is easy to check that

$$\sum_{l=1}^{\infty} p_i(l|A) = 1. \tag{11}$$

2) The Inter-group CTMC: The inter-group transitions can be modeled by an inter-group CTMC with 2^n states, each corresponding to one group of active BTSs A, and hence is referred to as a *lumped state* in view of the original CTMC. Two types of transitions can occur between the lumped states. First, the inter-group CTMC may transit from a lumped state B where queue i is empty to another lumped state A where queue i becomes nonempty, i.e.,

$$A = B \cup \{i\} \quad \text{and} \quad i \notin B. \tag{12}$$

The rate of such a transition is the rate that a packet arrives at queue i, λ_i. Second, the inter-group CTMC may transit from a lumped state A where queue i has length 1 to another lumped state B where queue i becomes empty (A and B satisfy (12)). The rate of such a transition is the probability that the queue length $l = 1$ times the service rate, and can be expressed as:

$$p_i(1|A) r_{i,A} = \left(1 - \frac{\lambda_i}{r_{i,A}}\right) r_{i,A} \tag{13}$$

$$= r_{i,A} - \lambda_i. \tag{14}$$

The transition rates that completely describe the inter-group CTMC of the lumped states are expressed as:

$$\hat{Q}(A, B) = \begin{cases} \lambda_i, & \text{if } B = A \cup \{i\}, i \notin A \\ r_{i,A} - \lambda_i, & \text{if } A = B \cup \{i\}, i \notin B \\ -\sum_{C: C \neq A} \hat{Q}(A, C), & \text{if } A = B \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

As an example, the CTMC of the lumped states is illustrated for the 2-BTS case in Fig. 4. The steady-state distribution $p(A)$ of the lumped states can be readily computed based on (15) using standard techniques in [21]. The key step therein is to invert a modified transition rate matrix.

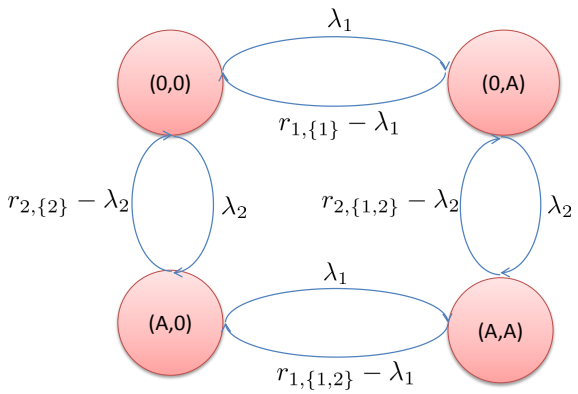

Fig. 4. The lumped CTMC for the 2-BTS interactive queueing model.

3) Average Delay: The average length of queue i in the modified CTMC can be calculated as follows:

$$\bar{l}_i = \sum_{\boldsymbol{l}} l_i p(\boldsymbol{l}) \tag{16}$$

$$= \sum_{\boldsymbol{l}} l_i p(A_{\boldsymbol{l}}) \prod_{j=1}^{n} p_j(l_j|A_{\boldsymbol{l}}) \tag{17}$$

$$= \sum_{A} p(A) \sum_{\boldsymbol{l}:A_{\boldsymbol{l}}=A} l_i \prod_{j=1}^{n} p_j(l_j|A) \tag{18}$$

$$= \sum_{A} p(A) \sum_{l_i:i\in A} l_i p_i(l_i|A) \left(\prod_{j\in A\setminus\{i\}} \sum_{l_j=1}^{\infty} p_j(l_j|A) \right) \tag{19}$$

$$= \sum_{A} p(A) \sum_{l_i:i\in A} l_i p_i(l_i|A) \tag{20}$$

$$= \sum_{A:i\in A} p(A) \sum_{l=1}^{\infty} l p_i(l|A) \tag{21}$$

$$= \sum_{A:i\in A} \frac{p(A) r_{i,A}}{r_{i,A} - \lambda_i}. \tag{22}$$

By Little's law, the average delay is given by $t_i = \bar{l}_i/\lambda_i$:

$$t_i = \sum_{A:i\in A} \frac{p(A) r_{i,A}}{(r_{i,A} - \lambda_i)\lambda_i}. \tag{23}$$

This is in general an approximation of the average delay of the original CTMC.

B. The Optimization Problem

Using the approximate average delay given by (23), we formulate the refined spectrum allocation problem as:

$$\underset{\boldsymbol{x},\boldsymbol{r},\boldsymbol{t}}{\text{minimize}} \quad \sum_{i=1}^{n} \frac{\lambda_i}{\sum_{j=1}^{n} \lambda_j} t_i \tag{P2a}$$

$$\text{subject to} \quad t_i = \sum_{A:i\in A} \frac{p(A) r_{i,A}}{(r_{i,A} - \lambda_i)\lambda_i}, \qquad \forall i \in N \tag{P2b}$$

$$r_{i,A} = \sum_{B\subset N} s_{i,B\cap A} x_B, \ \forall i \in N, \forall A \subset N \tag{P2c}$$

$$r_{i,N} > \lambda_i, \qquad \forall i \in N \tag{P2d}$$

$$x_B \geq 0, \qquad \forall B \subset N \tag{P2e}$$

$$\sum_{B\subset N} x_B = 1 \tag{P2f}$$

where $\boldsymbol{t} = [t_1, \ldots, t_n]$ and $\boldsymbol{r} = (r_{i,A})_{i\in N, A\subset N}$. Constraint (P2d) assures the stability of the CTMC within each lump state. It is not clear if Problem (P2) is convex due to the matrix inversion involved in calculating $p(A)$. Neither can we establish a counterpart to Theorem 1 in this case, i.e., there is no guarantee that the optimal solution uses at most n reuse patterns.

Nonetheless, we use a standard convex optimization algorithm to solve (P2). The simulations in Section V show that the algorithm always converges to the same solution regardless of the initial point. The resulting solution may divide the spectrum into a few more than n segments, but not by many.

C. Upper and Lower Bounds on the Delay

Several upper and lower bounds can be obtained for both the *actual* average delay under the refined model and its approximation (23). First, the conservative model's delay given by (4) is clearly an upper bound on the refined model's actual delay. Using tools developed in [28] for queues with time-varying link capacity, [17] developed bounds for queueing systems with coupled processors. The *first-degree bounds* therein are calculated assuming the best or worst service rates that decouple the queues. In fact, (4) corresponds exactly to the first-degree upper bound. The first-degree upper (lower) bound is usually loose in the light (heavy) traffic regime.

Tighter *second-degree bounds* are also presented in [17]. To lower (upper) bound the delay in cell i, the best (worst) rates are assumed to determine the utilization in interfering cells, i.e., the fraction of the time that those BTSs are transmitting. For example, to derive a second-degree upper bound for cell i, every other BTS j is assumed to transmit at its worst rate $r_{j,N}$. Denote $A_i(t)$ and $\bar{A}_i(t)$ as the set of active BTSs exclusive of BTS i at time t in the original interactive queuing system and under this assumption, respectively. It can be proved using a sample path argument that $A_i(t) \subset \bar{A}_i(t)$ at all time t. Thus the service rate at BTS i under this assumption is always lower than the service rate in the original interactive queuing system. This gives the delay upper bound. To derive the lower bound for cell i, all other BTSs are assumed to transmit at the best possible rates, $r_{j,\{j\}}$, $j \neq i$.

When assuming fixed rates at other BTSs, the queue at BTS i reduces to a queue with time-varying capacity. Processor sharing queues with time-varying capacity were studied in [28], which showed that the upper and lower bounds on residual work load in such queues can be obtained by considering the quasi-stationary regime and the fluid regime, respectively. In the quasi-stationary regime, the rest of the system evolves so slowly that BTS i only sees the initial states of the other BTSs. In the fluid regime the rest of the system evolves so quickly that BTS i only sees the average interference.

Under the worse-case transmit rate assumption, the probability that BTS j, $j \neq i$ transmits is:

$$\bar{p}_j = \frac{\lambda_j}{r_{j,N}}. \tag{24}$$

The probability that the other $n-1$ BTSs are in state $\bar{A}_i$ is:

$$\bar{\pi}_i(\bar{A}_i) = \prod_{j \in \bar{A}_i} \bar{p}_j \prod_{l \notin \bar{A}_i, l \neq i} (1 - \bar{p}_l). \tag{25}$$

The second-degree upper bound is finally given by taking the expectation over the distribution of all possible states $\bar{A}_i$:

$$\bar{t}_i = \sum_{\bar{A}_i \subset (N \setminus \{i\})} \bar{\pi}_i(\bar{A}_i) \frac{1}{r_{i,\bar{A}_i \cup \{i\}} - \lambda_i}. \tag{26}$$

Under the best possible rate assumption, the probability of BTS j, $j \neq i$ being active is:

$$\underline{p}_j = \frac{\lambda_j}{r_{j,\{j\}}}. \tag{27}$$

The corresponding probability of state $\underline{\mathcal{A}}_i$ for the $n-1$ BTSs except BTS i is:

$$\underline{\pi}_i(\underline{A}_i) = \prod_{j \in \underline{A}_i} \underline{p}_j \prod_{l \notin \underline{A}_i, l \neq i} (1 - \underline{p}_l). \tag{28}$$

The second-degree lower bound is calculated using the average rate:

$$\underline{t}_i = \frac{1}{\sum_{\underline{A}_i \subset (N \setminus \{i\})} \underline{\pi}_i(\underline{A}_i) r_{i,\underline{A}_i \cup \{i\}} - \lambda_i}. \tag{29}$$

The accuracy of the refined approximation for the average delay developed in Section IV-A is guaranteed by the following theorem:

Theorem 2: In a n-BTS interactive queueing system, the refined approximate mean packet sojourn time provided by (23) is between the second-degree upper and lower bounds in (26) and (29), i.e., $\underline{t}_i < t_i < \bar{t}_i$.

The proof of *Theorem* 2 is given in Appendix A.

D. Throughput Optimality of the Proposed Schemes

The throughput region of an interactive queueing system is given in [30] in the context of resource allocation through coordinated scheduling in the time domain. Adapting the result to the current setting gives the throughput region of the n-BTS network with given spectral efficiencies $(s_{i,B})_{i \in N,\ B \subset N}$ as:

$$\Lambda = \left\{ (\lambda_1, \ldots, \lambda_n) \,\middle|\, \exists\, (x_B)_{B \subset N}, \text{ s.t. } r_i \geq \lambda_i, \forall i \in N \right\} \tag{30}$$

where r_i is defined as in (2). That is, for any rate tuple in the interior of Λ, there exists a spectrum allocation that stabilizes the interactive queueing system, whereas for any rate tuples outside the region, there is no allocation that can stabilize all the queues.

Theorem 3: The proposed conservative and refined spectrum allocation schemes are both throughput optimal, namely, they both achieve the entire throughput region as given in (30).

Proof: Problems (P1) and (P2) both include the constraint, $r_i > \lambda_i\ \forall i \in N$. Hence the feasible regions for the two proposed schemes are exactly the throughput region given by (30). Hence their throughput optimality follows. ■

The result in [30] only guarantees that for any rate tuple that can be stabilized, there exists a spectrum allocation $r_i \geq \lambda_i,\ \forall i \in N$. However, there might be another spectrum allocations that also stabilizes the system, but have $r_i < \lambda_i$ at some of the BTSs. The intuition is that although the worst case service rate is less than the traffic arrival rate at some queues, those queues will also operate at other states A with less interference and higher rates. By Loynes theorem [31], the interactive queueing system will be stable, as long as the average service rate exceeds the packet arrival rate at each queue. Determining the stability of an interactive queueing system with fixed arrival rates $(\lambda_i)_{i \in N}$ and service rates $(r_{i,B})_{i \in N,\ B \subset N}$ is more complicated. The result for two interactive queues is presented in [32].

V. Numerical Results

Throughout this section, the noise PSD is 0.125×10^{-6} μW/Hz, the average packet size is $L = 1$ Mbit and the total bandwidth is $W = 20$ MHz in all simulations. In the homogenous setup, the path-loss exponent is 3, and the transmit PSD is 1 μW/Hz for all BTSs. In the heterogenous setup, the transmit PSD and the pathloss exponent of the macro BTS are 10 μW/Hz and 2.8, respectively; while the transmit PSD and the pathloss exponent for each pico BTS are 1 micro-watt/Hz and 3.4, respectively. The traffic distribution at each BTS is assumed to be proportional to its worst-case transmit rate, i.e., $\lambda_i \propto r_{i,N},\ i \in N$. When the average load increases, $\lambda_1, \cdots, \lambda_N$ increase proportionally. However, the key conclusions apply to an arbitrary traffic distribution. The results for actual packet delays with service rates given by (3), are obtained by simulating the equivalent discrete time Markov chain (DTMC) of the CTMC (7) using the uniformization method [33] for 10^5 time intervals.

A. One-Dimensional Example

We first present the one-dimensional example in Fig. 5 to demonstrate that the proposed spectrum allocation scheme is topology- and traffic-aware. Seven BTSs, represented by the triangles, are randomly placed on a line segment. The relative traffic load at each BTS is depicted by the height of each rectangle below it, whereas the spectrum allocated using the refined scheme is depicted using the rectangle above it. A solid block depicts an active reuse pattern. Evidently, both spatial reuse and local orthogonalization are accomplished by the allocation. For example, the right most cell has much heavier

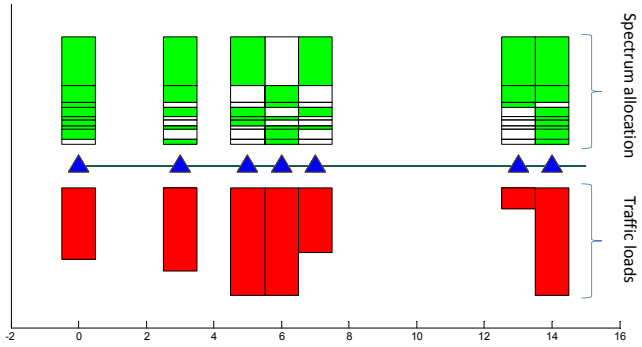

Fig. 5. One dimensional refined spectrum allocation.

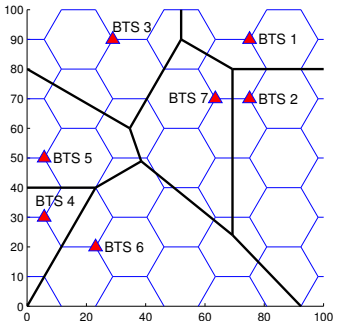

Fig. 6. The topology of the 7-pico BTS network.

traffic than its immediate neighbor, and is thus allocated much more spectrum. Also, the spectrum allocated to the middle BTS is mostly orthogonal with those of the two close neighbors.

B. Two-Dimensional Simulation Model

To illustrate the performance of the proposed conservative and refined schemes, we adopt the quantized HetNet model in [34]. A 100×100 m^2 area is partitioned into hexagons, where the distance between nearest centers is 20 m. In the simulation, 7 BTSs are uniformly randomly dropped at the vertices of the hexagons. Each UE location is approximated by the center of its hexagon, and is assigned to its nearest BTS. The channel gain are determined using the standard path-loss model. The topology of the network is shown in Fig. 6, where the locations of the BTSs are denoted by the triangles. The Voronoi region of each pico cell is also shown. The average spectral efficiency of BTS i is calculated as the mean of the spectral efficiencies of the hexagons it serves.

C. Delay Comparison

The delays using the refined and conservative approximations are compared with the second-degree upper and lower bounds in Fig. 7. The average delay versus average traffic arrival rate curves in Fig. 7 are based on the same spectrum allocation. The conservative approximation, also known as the first-degree upper bound, is coarser than the second-degree upper bound. As predicted by Theorem 2, the refined approximation is between the second-degree upper and lower

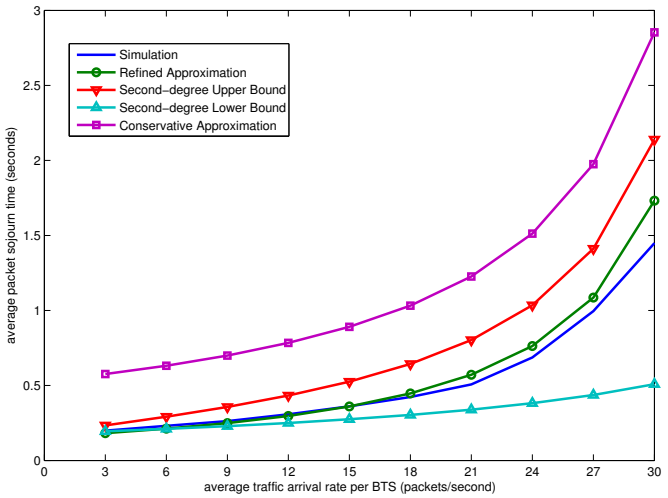

Fig. 7. Comparison of the conservative and refined approximations with the second-degree upper and lower bounds.

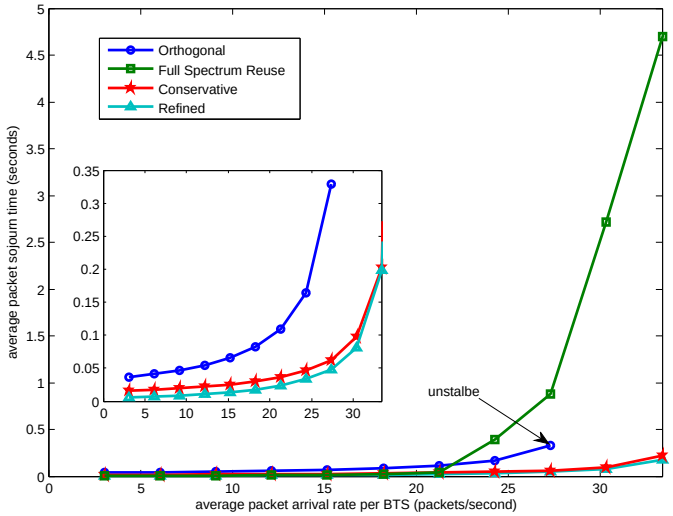

Fig. 8. Comparison of the conservative and refined allocations with the orthogonal and full-reuse allocations

bounds. In addition, the refined approximation is also quite accurate within the feasible region.

We compare the conservative and refined spectrum allocations with simpler orthogonal and full-reuse allocations for different traffic loads in Fig. 8. The orthogonal allocation is the optimal one among all feasible orthogonal allocations. The figure shows that the orthogonal allocation becomes unstable after the average packet arrival rate reaches 27 packets/second, suggesting this is the throughput region under the orthogonal allocation. The delay under the full-reuse allocation grows much faster after the average packet arrival rate rises above 21 packets/second. A general observation is that as the difference in spectral efficiency under the orthogonal and full-reuse allocations grows larger, the delay under full-reuse also increases more rapidly with the traffic. As suggested by the numerical results presented in Section V-E, the full-reuse allocation also

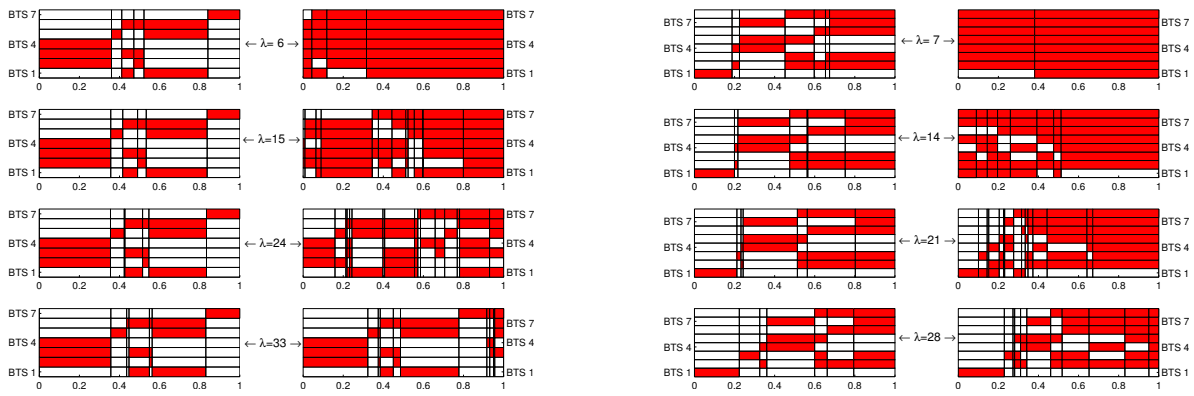

(a) The homogeneous setup with 7 pico BTSs.

(b) The heterogeneous setup with 1 macro BTS and 7 pico BTSs

Fig. 9. Optimal conservative (left) and refined (right) allocations with different packet arrival rates λ.

becomes unstable at 27 pacekts/second.[2] Hence, the proposed conservative and refined allocations achieve a larger throughput region than the other schemes. We can also observe from Fig. 8 that the proposed conservative and refined allocations achieve significant gains in the heavy-traffic regime.

Since the delay under the full-reuse allocation becomes large in the heavy-traffic regime, the orthogonal, conservative and refined allocations are compared separately in the inlet in Fig. 8. Since the average sojourn time under the orthogonal allocation can be exactly calculated from (P1a), the minimum of (P1) given by the optimal conservative allocation is always no greater than the actual delay of the orthogonal allocation. Moreover, since the conservative approximation is an upper bound on the actual delay, the orthogonal allocation is always worse than the conservative allocation as shown in the inlet. The refined allocation always outperforms the conservative allocation due to the more accurate approximation of the actual delay. In the light-traffic regime, the refined allocation reduces the average delay by about 60% compared to the conservative allocation; while both provide significant delay reduction compared to the orthogonal allocation. The advantage of using the refined allocation over the conservative allocation decreases as the traffic increases. In the heavy-traffic regime, the difference is negligible.

The optimal conservative and refined allocations are shown in Fig. 9a for different traffic loads. The widths of the rectangles represent the bandwidths of the active reuse patterns. The solid ones in each row are the spectrum segments that are used by the corresponding BTS. The number of active reuse patterns employed by the conservative scheme is exactly the same as the number of cells (c.f., Theorem 1). In the light-traffic regime, the refined allocation is close to full-reuse, and as the traffic increases it becomes closer to the conservative allocation. This is because in the light-traffic regime all BTSs are inactive most of the time, whereas in the heavy-traffic regime, all queues are mostly occupied, which perform similarly to n independent M/M/1 queues with the worst-case service rates.

[2]The finite delay of the full-reuse allocation at arrival rate of 27 packets/second and even higher is due to finite simulation time.

D. Heterogenous Setup

Fig. 9b illustrates the conservative and refined allocations for a HetNet with a single macro (large) cell and several pico (small) cells. The simulation setup is the same as described in Section V-B, with an additional macro BTS added at the center of the network. The proposed conservative and refined allocations are compared with the simple optimal orthogonal and full-reuse allocations at different traffic levels. The delay performance versus different loads looks similar to the one shown in Fig. 8, and is omitted.

As shown in Fig. 9b, the conservative allocation orthogonalizes the spectrum use between the macro BTS and the pico BTSs in all traffic regimes. (BTS 1 is the macro BTS, whose spectrum allocation is shown at the bottom of each subplot.) This is because the macro BTS causes significant interference to the pico BTSs under the backlogged interference assumption. In the refined allocation, the macro BTS can still share part of the spectrum with the pico BTSs in the light-traffic regime. As the traffic load increases, an orthogonal spectrum allocation is observed between macro and pico tiers. Of course, the proposed spectrum allocation schemes can be applied to other more general HetNets, as long as the spectral efficiencies under different sharing combinations along with the traffic arrival rates can be obtained by a central controller.

E. Utilization

We present the utilization ratio of each BTS under different traffic loads in Fig. 10. In Fig. 10a, where the traffic is very light, the refined allocation has very similar utilization as full-reuse. The utilization ratios are much higher under the conservative allocation due to the worst-case (smallest) transmit rate assumption. Since the BTSs are active less than 10% of the time, the orthogonal allocation is highly inefficient. Fig. 10b presents the utilization ratios at an arrival rate where the full-reuse allocation starts to incur high delay as shown in Fig. 8. The conservative and refined allocations have much lower BTS utilization ratios compared to the

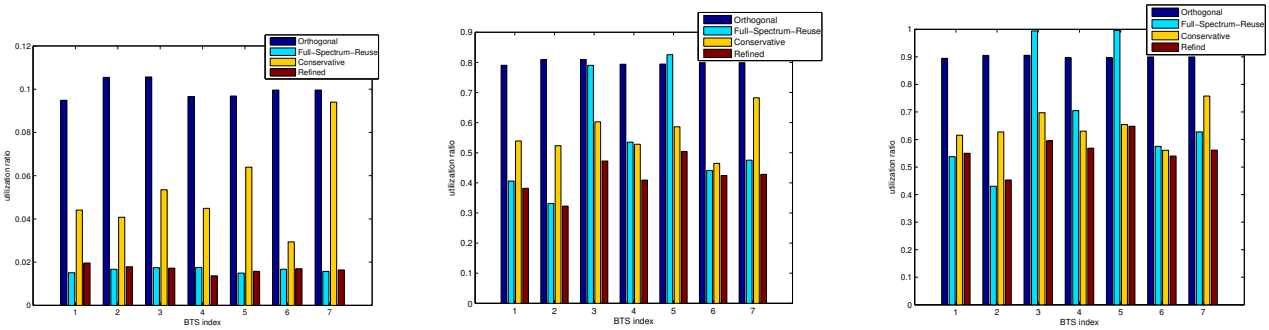

(a) average packet arrival rate = 3 (b) average packet arrival rate = 24 (c) average packet arrival rate = 27

Fig. 10. BTS utilization ratios for different average packet arrival rates.

full-reuse allocation, especially at BTS 3 and BTS 5. As the traffic demand increases, as shown in Fig. 10c, BTSs 3 and 5 become saturated under full-reuse, which also means the system becomes unstable. The conservative and refined allocations still maintain the stability of the queues and have much lower utilization ratios than the orthogonal allocation. As suggested in Fig. 8, if the traffic increases further to 30 packets/second, the orthogonal allocation also becomes unstable.

F. Delay Distribution

The CDF of the number of packets in all queues for the entire network is shown in Fig. 11. In the light-traffic regime, as shown in Fig. 11a, the refined and full-reuse allocations have the smallest number of packets at all percentiles. The conservative allocation has more packets than those two, but fewer packets than the orthogonal allocation. The queues are empty 90% of the time under all four allocations. At average traffic arrival rate of 24 packets/second, the refined and conservative allocations both have much fewer packets at all percentiles than the other two, as shown in Fig. 11b. The orthogonal allocation has a more balanced queue length compared with full-reuse due to fixed service rates. At even higher traffic loads, as shown in Fig. 11c, the full-reuse allocation becomes unstable. The conservative and refined allocations have substantial advantages over the other two in this heavy-traffic regime.

G. Power Control

The discussions so far have been based on the important assumption that the spectral efficiencies when shared by different combinations of BTSs, i.e., $s_i(A),\ \forall i \in N,\ A \subset N$, are fixed. This assumption enables us to simplify the relation between spectrum allocation and flow level service rate. In fact, the service rate is a linear function of the allocated bandwidths under this fixed spectral efficiency assumption. This assumption is valid if all BTSs transmit with fixed power spectral density.

However, in practice we may have a fixed total transmit power constraint at each BTS. Power control is not present in the current formulation. For example, the spectral efficiencies for the orthogonal allocation should be higher then shown,

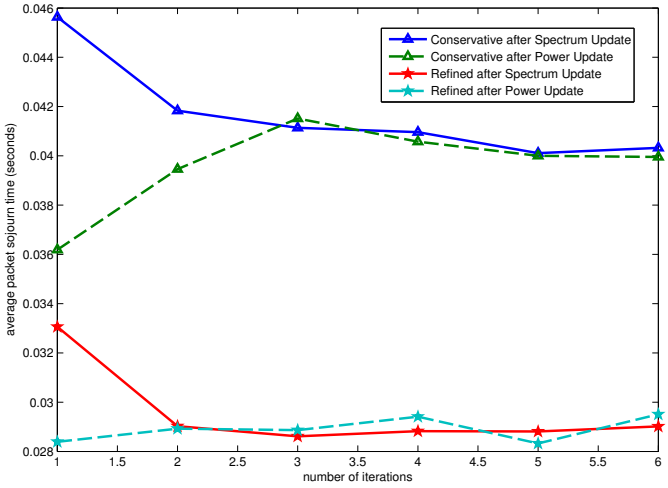
Conservative after Spectrum Update
Conservative after Power Update
Refined after Spectrum Update
Refined after Power Update
average packet sojourn time (seconds)
number of iterations

Fig. 12. Average packet sojourn time versus number of iterations.

since each BTS would concentrate all the transmit power on its exclusive spectrum. The joint power and spectrum allocation problem in its general form can be formulated as optimizing a continuous power spectral density function at each BTS subject to a total transmit power constraint. The discretized version is proved to be NP-hard in [35]. The continuous version is also difficult.

We next take a simplified approach by alternatively updating the spectrum and power allocations. At the beginning, the spectral efficiencies are fixed assuming each BTS uniformly allocates its maximum transmit power across the entire spectrum. Then, we iterate the following steps:

1) Update the spectrum allocation $x_B,\ \forall B \subset N$ with the current $s_{i,A},\ \forall i \in N,\ A \subset N$ by solving the proposed spectrum allocation problems.
2) Update the spectral efficiencies $s_{i,A},\ \forall i \in N,\ A \subset N$ with the current $x_B,\ \forall B \subset N$, by letting each BTS uniformly allocate its maximum transmit power over the spectrum segments assigned to it.

The iterations continue until the spectrum allocation $\boldsymbol{x}$ converges. (This is not guaranteed.) The average packet sojourn time after the spectrum allocation update and the spectral efficiency update at each iteration is shown in Fig. 12 for an average packet arrival rate per BTS of 24 packets/second.

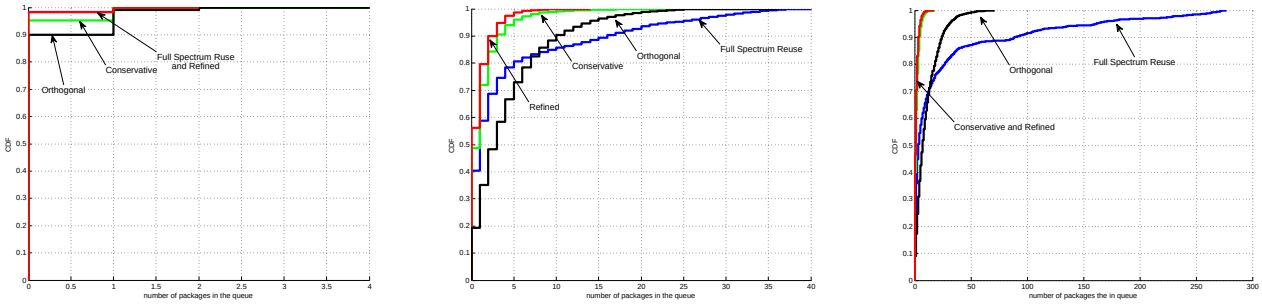

(a) average packet arrival rate = 3 (b) average packet arrival rate = 24 (c) average packet arrival rate =27

Fig. 11. CDF's of number of packets in each queue for different average packet arrival rates.

The figure shows that the delay performance converges very quickly for both conservative and refined allocations. The mean sojourn times for both allocations decrease substantially after the first spectral efficiency update. This is because at this average packet arrival rate, both allocations orthogonalize the spectrum use among neighboring BTSs to some extent. (The small variations in these curves are due to the limited simulation time.) This kind of convergence behavior can be expected in general, except for networks with very few BTSs. This is because spatial reuse will occur in the conservative and refined allocations for relatively large networks even in the heavy-traffic regime. Since each BTS will use a fairly large amount of the spectrum, the spectral efficiencies will not change much after several iterations.

VI. Conclusion

Traffic-driven spectrum allocation performed over a slow timescale in a HetNet has been studied. Two efficient schemes based on optimization have been developed and shown to be highly effective. The proposed model and queueing analysis for densely deployed HetNets depart from the traditional single cell and regular hexagonal cellular network models. Also, the solution fully adapts to the topology of the network, rather than attempting to average over random realization using stochastic geometry.

The conservative and refined allocations significantly reduce the average delay by exploiting the network topology and the different loads across BTSs. The problem formulation and results can be generalized to arbitrary link performance measure and convex objective function, beyond the specific choice of Shannon spectral efficiency and queueing delay. In particular, the formulation directly applies to the sum rate and weighted sum rate objectives.

An important extension of the work is to develop a general framework for radio resource managment by incorporating user association, multiple antennas, and multiple radio access technologies (RATs). The effort is well-matched to recent proposals for a cloud radio access network (C-RAN). Ultimately, the goal is to also extended the framework to address large networks with hundreds to thousands of cells, thus pushing joint radio resource management to an unprecedented scale.

Appendix A: Proof of Theorem 2

First we will prove $t_i < \bar{t}_i$. Both (23) and (26) can be regarded as weighted sum of $\frac{r_{i,A}}{(r_{i,A}-\lambda_i)\lambda_i}$. The weight is the probability of the system being in state A under different approximation. Based on the worst-case rate assumption for deriving $\bar{t}_i$, we can form a similar lumped chain as (15) with the transition rate matrix:

$$\bar{Q}(A,B) = \begin{cases} \lambda_i, & \text{if } B = A \cup \{i\}, i \notin A \\ (r_{i,N} - \lambda_i), & \text{if } A = B \cup \{i\}, i \notin B \\ -\sum_{C \neq A} \bar{Q}(A,C), & \text{if } A = B \\ 0, & \text{otherwise.} \end{cases} \tag{31}$$

The lumped chains (15) and (31) can be interpreted as the CTMC of the following Markov process: There are n interactive queues each with capacity 1. Packet arrivals at queue i follow a Poisson process with rate λ_i. The size of each packet is independently exponentially distributed with unit mean. The service rates are given by the corresponding rate matrices of (15) and (31), which are state dependent. If a packet arrives at an empty queue, it is immediately served. Otherwise, the packet in service will be discarded and the newly arrived packet immediately starts being served. Denote S and $\bar{S}$ respectively as the corresponding systems with probability transition matrix given by (15) and (31). Let $A(t)$ and $\bar{A}(t)$ be the sets of active queues at time t. Denote the residual loads (in bits) by $\boldsymbol{l}(t)$ and $\bar{\boldsymbol{l}}(t)$, where $\boldsymbol{l}(t),\ \bar{\boldsymbol{l}}(t) \in \mathbb{R}_+^n$ and the ith elements $l_i(t)$ and $\bar{l}_i(t)$ are the residual loads in queue i of the two systems. Using a sample path argument, we can show:

$$\begin{aligned} A(t) &\subset \bar{A}(t) \\ \boldsymbol{l}(t) &\leq \bar{\boldsymbol{l}}(t) \end{aligned} \tag{32}$$

at any time instance t. The inequality $\boldsymbol{l}(t) \leq \bar{\boldsymbol{l}}(t)$ is element-wise, i.e., $l_i(t) \leq \bar{l}_i(t),\ i = 1, \ldots, n$.

Assuming both systems evolve under the same packet arrival realization, we can prove (32) by induction. At $t = 0$, $A(0) = \bar{A}(0) = \emptyset$ and $\boldsymbol{l}(0) = \bar{\boldsymbol{l}}(0) = 0$. Assume (32) is true at $t = \tau > 0$. Then, (32) still holds before any arrival or departure happens at $t = \tau + \delta$. This is because $A(t) \subset \bar{A}(t),\ \forall t \in [\tau, \tau + \delta)$ implies the service rate of any queue i in S is larger than the service rate of queue i in $\bar{S}$ within this time interval. At $t = \tau + \delta$, we have two cases.

Case 1. An arrival happens at queue j. Since both systems follows the same sample path of the arrivals, we still have $A(\tau+\delta) \subset A(\tau+\delta)$. The residual load at other queues does not change from $\tau+\delta_-$ to $\tau+\delta$. At queue j, $\boldsymbol{l}_j(\tau+\delta) = \bar{\boldsymbol{l}}_j(\tau+\delta)$.
Case 2. A departure happens at queue j. Since $\boldsymbol{l}(t) \leq \bar{\boldsymbol{l}}(t),\ \forall t \in [\tau, \tau+\delta)$ and the service rate at each queue in S is no less than the service rate at the same queue in $\bar{S}$, this departure must be in S (It could be the case that a departure also happens at queue j in $\bar{S}$ at $t=\tau+\delta$.). It immediately implies (32) at $t=\tau+\delta$.

Because the Markov processes in both systems are ergodic, we have:

$$\begin{aligned} t_i &= \lim_{T\to\infty} \frac{1}{T} \int_0^T f_i\left(A(t)\right) dA(t) \\ \bar{t}_i &= \lim_{T\to\infty} \frac{1}{T} \int_O^T f_i\left(\bar{A}(t)\right) d\bar{A}(t), \end{aligned} \tag{33}$$

where the function $f_i(A)$ is:

$$f_i(A) = \begin{cases} \dfrac{r_{i,A}}{(r_{i,A}-\lambda_i)\lambda_i} & \text{if } i \in A \\ 0 & \text{otherwise} \end{cases}$$

It is easy to see that if $A \subset B$ then $f_i(A) \leq f_i(B)$. Applying this and (32) to (33), we can obtain $t_i < \bar{t}_i$, where the strict inequality is due to ergodicity.

To prove $t_i > \underline{t}_i$, we first introduce an intermediate variable t_i':

$$t_i' = \frac{1}{\sum_{A: i\in A} P(A) \frac{r_{i,A}}{\lambda_i} r_{i,A} - \lambda_i}. \tag{34}$$

According to (15), it is easy to check $\sum_{A:i\in A} P(A)\frac{r_{i,A}}{\lambda_i} = 1$. By Jensen's inequality, we immediately have $t_i \geq t_i'$. Hence, we only need to show $\hat{\underline{t}}_i > \underline{t}_i$. Again we can form a lumped chain with the highest service rate assumption, whose transition rate matrix is:

$$\underline{Q}(A,B) = \begin{cases} \lambda_i, & \text{if } B = A \cup \{i\}, i \notin A \\ \left(r_{i,\{i\}} - \lambda_i\right), & \text{if } A = B \cup \{i\}, i \notin B \\ -\sum_{C: C\neq A} \underline{Q}(A,C), & \text{if } A = B \\ 0, & \text{otherwise.} \end{cases} \tag{35}$$

Using a similar sample path argument we can prove the expected service rate in (34) is less than the expected service rate in (29), which directly implies $t_i' > \underline{t}_i$.

REFERENCES

[1] B. Zhuang, D. Guo, and M. L. Honig, "Traffic driven resource allocation in heterogenous wireless networks," in *Proc. IEEE GLOBECOM*, pp. 1546–1551, Dec 2014.

[2] H. Lei, L. Zhang, X. Zhang, and D. Yang, "A novel multi-cell OFDMA system structure using fractional frequency reuse," in *Proc. IEEE PIMRC*, pp. 1–5, Sept 2007.

[3] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, "What will 5G be?," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, 2014.

[4] J. Liu, W. Xiao, and A. C. K. Soong, *Design and Deployment of Small Cell Networks*, ch. Dense Networks of Small Cells. Cambridge University Press, 2014.

[5] A. L. Stolyar and H. Viswanathan, "Self-organizing dynamic fractional frequency reuse in OFDMA systems," in *Proc. IEEE INFOCOM*, Apr. 2008.

[6] R. Chang, Z. Tao, J. Zhang, and C.-C. Kuo, "Multicell OFDMA downlink resource allocation using a graphic framework," *IEEE Trans. Veh. Technol.*, vol. 58, pp. 3494–3507, Sept 2009.

[7] S. H. Ali and V. C. M. Leung, "Dynamic frequency allocation in fractional frequency reused OFDMA networks," *IEEE Trans. Wireless Commun.*, vol. 8, pp. 4286–4295, Aug. 2009.

[8] R. Madan, J. Borran, A. Sampath, N. Bhushan, A. Khandekar, and T. Ji, "Cell association and interference coordination in heterogeneous LTE-A cellular networks," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 9, pp. 1479–1489, 2010.

[9] W.-C. Liao, M. Hong, Y.-F. Liu, and Z.-Q. Luo, "Base station activation and linear transceiver design for optimal resource management in heterogeneous networks," *IEEE Trans. Signal Process.*, vol. 62, no. 15, pp. 3939–3952, 2014.

[10] J. Huang, V. Subramanian, R. Agrawal, and R. Berry, "Joint scheduling and resource allocation in uplink OFDM systems for broadband wireless access networks," *IEEE J. Sel. Areas Commun.*, vol. 27, pp. 226–234, February 2009.

[11] D. Fooladivanda and C. Rosenberg, "Joint resource allocation and user association for heterogeneous wireless cellular networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 1, pp. 248–257, 2013.

[12] G. Lim, C. Xiong, L. J. Cimini, and G. Y. Li, "Energy-efficient resource allocation for OFDMA-based multi-RAT networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 5, pp. 2696–2705, 2014.

[13] K. Shen and W. Yu, "Distributed pricing-based user association for downlink heterogeneous cellular networks," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1100–1113, 2014.

[14] S. Deb, P. Monogioudis, J. Miernik, and J. P. Seymour, "Algorithms for enhanced inter-cell interference coordination (eICIC) in LTE HetNets," *IEEE/ACM Trans. Netw.*, vol. 22, no. 1, pp. 137–150, 2014.

[15] H. Dhillon, R. K. Ganti, and J. G. Andrews, "Load-aware modeling and analysis of heterogeneous cellular networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 4, pp. 1666–1677, 2013.

[16] A. Shojaeifard, K. A. Hamdi, E. Alsusa, D. K. C. So, and J. Tang, "A unified model for the design and analysis of spatially-correlated load-aware HetNets," *IEEE Trans. Commun.*, vol. 62, no. 11, pp. 4110–4125, 2014.

[17] T. Bonald, S. Borst, N. Hegde, and A. Proutiére, "Wireless data performance in multi-cell scenarios," in *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS/Performance, (New York, NY, USA), pp. 378–380, ACM, 2004.

[18] B. Rengarajan and G. de Veciana, "Architecture and abstractions for environment and traffic aware system-level coordination of wireless networks: the downlink case," in *Proc. IEEE INFOCOM*, 2008.

[19] Q. Kuang, W. Utschick, and A. Dotzler, "Optimal joint user association and resource allocation in heterogeneous networks via sparsity pursuit," *http://arxiv.org/abs/1408.5091*, 2014.

[20] M. Haviv, *Queues: A Course in Queueing Theory*, vol. 191. Springer, 2013.

[21] R. Nelson, *Probability, stochastic processes, and queueing theory: the mathematics of computer performance modeling*. Springer, 1995.

[22] D. P. Bertsekas, *Nonlinear programming*. Athena Scientific, 1999.

[23] H. G. Eggleston, *Convexity*. No. 47, Cambridge University Press Archive, 1958.

[24] D. Bertsimas and J. N. Tsitsiklis, *Introduction to linear optimization*, vol. 6. Athena Scientific Belmont, MA, 1997.

[25] G. Fayolle and R. Iasnogorodski, "Two coupled processors: the reduction to a riemann-hilbert problem," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 47, no. 3, pp. 325–351, 1979.

[26] I. J. B. F. Adan, J. Wessels, and W. H. M. Zijm, "A compensation approach for two-dimensional markov processes," *Advances in Applied Probability*, vol. 25, pp. 783–817, Dec. 1993.

[27] J. W. Cohen and O. J. Boxma, *Boundary value problems in queueing system analysis*. Elsevier, 2000.

[28] F. Delcoigne, A. Proutire, and G. Rgni, "Modeling integration of streaming and data traffic," *Performance Evaluation*, vol. 55, no. 3-4, pp. 185–209, 2004.

[29] B. Zhuang, D. Guo, and M. L. Honig, "Traffic-driven spectrum allocation in heterogeneous networks," *http://arxiv.org/abs/1408.6011*, 2014.

[30] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, pp. 1936–1948, Dec. 1992.

[31] R. M. Loynes, "The stability of a queue with non-independent inter-arrival and service times," in *Cambridge Philos. Soc*, vol. 58, pp. 497–520, Cambridge Univ Press, 1962.
[32] C. Knessl and J. A. Morrison, "Heavy traffic analysis of two coupled processors," *Queueing Systems*, vol. 43, no. 3, pp. 173–220, 2003.
[33] W. J. Stewart, *Probability, Markov chains, queues, and simulation: the mathematical basis of performance modeling*. Princeton University Press, 2009.
[34] B. Zhuang, M. L. Honig, and D. Guo, "Energy management of dense wireless heterogeneous networks over slow timescales," in *Proc. Allerton Conf. Commun., Control, & Computing*, 2012.
[35] Z.-Q. Luo and S. Zhang, "Dynamic spectrum management: Complexity and duality," *IEEE J. Sel. Topics Signal Process.*, vol. 2, pp. 57–73, Feb 2008.

Binnan Zhuang received his B.S. degree from Electronic Engineering Department of Tsinghua University in July 2009 the M.S. degree EECS department of Northwestern University. He is pursuing his Ph.D. degree at EECS department of Northwestern University. His research interests include wireless communications, communication network and network optimization. His current research focus on interference and resource management in heterogeneous networks.

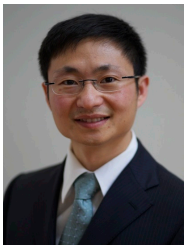

Dongning Guo (S'97-M'05-SM'11) received the B.Eng. degree from the University of Science and Technology of China, Hefei, China, the M.Eng. degree from the National University of Singapore, Singapore, and the M.A. and Ph.D. degrees from Princeton University, Princeton, NJ, USA. In 2004, he joined the faculty of Northwestern University, Evanston, IL, USA, where he is currently an Associate Professor in the Department of Electrical Engineering and Computer Science. From 1998 to 1999, he was an R&D Engineer in the Center for Wireless Communications, Singapore. He has been an Associate Editor of IEEE Transactions on Information Theory and a Guest Editor of a Special Issue of IEEE Journal on Selected Areas in Communications. He is an Editor of Foundations and Trends in Communications and Information Theory. Dr. Guo received the Huber and Suhner Best Student Paper Award in the International Zurich Seminar on Broadband Communications in 2000 and is the corecipient of the IEEE Marconi Prize Paper Award in Wireless Communications in 2010. He is also the recipient of the National Science Foundation Faculty Early Career Development (CAREER) Award in 2007. His research interests include information theory, communications, and networking.

Michael L. Honig (S'80-M'81-SM'92-F'97) received the B.S. degree in electrical engineering from Stanford University in 1977, and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley, in 1978 and 1981, respectively. He subsequently joined Bell Laboratories in Holmdel, NJ, where he worked on local area networks and voiceband data transmission. In 1983 he joined the Systems Principles Research Division at Bellcore, where he worked on Digital Subscriber Lines and wireless communications. Since the Fall of 1994, he has been with Northwestern University where he is a Professor in the Department of Electrical and Computer Engineering. He has held visiting scholar positions at the Naval Research Laboratory (San Diego), the University of California, Berkeley, the University of Sydney, Princeton University, and the Technical University of Munich. He has also worked as a freelance trombonist.

Dr. Honig has served as an Editor for the IEEE Transactions on Information Theory and the IEEE Transactions on Communications, and as Guest Editor for several journals. He has also served as a member of the Board of Governors for the Information Theory Society. He is the recipient of a Humboldt Research Award for Senior U.S. Scientists, and the corecipient of the 2002 IEEE Communications Society and Information Theory Society Joint Paper Award and the 2010 IEEE Marconi Prize Paper Award.